

CMSC 754: Lecture 1

Introduction to Computational Geometry

What is Computational Geometry? “Computational geometry” is a term claimed by a number of different groups. The term was coined perhaps first by Marvin Minsky in his book “Perceptrons”, which was about pattern recognition, and it has also been used often to describe algorithms for manipulating curves and surfaces in solid modeling. It’s most widely recognized use, however, is to describe the subfield of algorithm theory that involves the design and analysis of efficient algorithms for problems involving geometric input and output.

The field of computational geometry grew rapidly in the late 70’s and through the 80’s and 90’s, and it is still a very active field of research. Historically, computational geometry developed as a generalization of the study of algorithms for sorting and searching in 1-dimensional space to problems involving multi-dimensional inputs. Because of its history, the field of computational geometry has focused mostly on problems in 2-dimensional space and to a lesser extent in 3-dimensional space. When problems are considered in multi-dimensional spaces, it is often assumed that the dimension of the space is a small constant (say, 10 or lower). Nonetheless, recent work in this area has considered a limited set of problems in very high dimensional spaces, particularly with respect to approximation algorithms. In this course, our focus will be largely on problems in 2-dimensional space, with occasional forays into spaces of higher dimensions.

Because the field was developed by researchers whose training was in discrete algorithms (as opposed to more continuous areas such as a numerical analysis or differential geometry) the field has also focused principally on the *discrete* aspects of geometric problem solving. The mixture of discrete and geometric elements gives rise to an abundance of interesting questions. (For example, given a collection of n points in the plane, how many pairs of points can there be that are exactly one unit distance apart from each other?) Another distinctive feature is that computational geometry primarily deals with straight or flat objects (lines, line segments, polygons, planes, and polyhedra) or simple curved objects such as circles. This is in contrast, say, to fields such as solid modeling, which focus on issues involving curves and surfaces and their representations.

Computational geometry finds applications in numerous areas of science and engineering. These include computer graphics, computer vision and image processing, robotics, computer-aided design and manufacturing, computational fluid-dynamics, and geographic information systems, to name a few. One of the goals of computational geometry is to provide the *basic geometric tools* needed from which application areas can then build algorithms and *theoretical analytic tools* needed to analyze the performance of these algorithms. There has been significant progress made towards this goal, but it is still far from being fully realized.

A Typical Problem in Computational Geometry: Here is an example of a typical problem, called the *shortest path problem*. Given a set polygonal obstacles in the plane, find the shortest obstacle-avoiding path from some given start point to a given goal point (see Fig. 1). Although it is possible to reduce this to a shortest path problem on a graph (called the *visibility graph*, which we will discuss later this semester), and then apply a nongeometric algorithm such as Dijkstra’s algorithm, it seems that by solving the problem in its geometric domain it should

be possible to devise more efficient solutions. This is one of the main reasons for the growth of interest in geometric algorithms.




Fig. 1: Shortest path problem.

The measure of the quality of an algorithm in computational geometry has traditionally been its *asymptotic worst-case running time*. Thus, an algorithm running in $O(n)$ time is better than one running in $O(n \log n)$ time which is better than one running in $O(n^2)$ time. (This particular problem can be solved in $O(n^2 \log n)$ time by a fairly simple algorithm, in $O(n \log n)$ by a relatively complex algorithm, and it can be approximated quite well by an algorithm whose running time is $O(n \log n)$.) In some cases *average case* running time is considered instead. However, for many types of geometric inputs (this one for example) it is difficult to define input distributions that are both easy to analyze and representative of typical inputs.

Overview of the Semester: Here are some of the topics that we will discuss this semester.

Convex Hulls: Convexity is a very important geometric property. A geometric set is *convex* if for every two points in the set, the line segment joining them is also in the set. One of the first problems identified in the field of computational geometry is that of computing the smallest convex shape, called the *convex hull*, that encloses a set of points (see Fig. 2).




Fig. 2: Convex hulls, intersections, and polygon triangulation.


Intersections: One of the most basic geometric problems is that of determining when two sets of objects intersect one another. Determining whether complex objects intersect often reduces to determining which individual pairs of primitive entities (e.g., line segments) intersect (see Fig. 2). We will discuss efficient algorithms for computing the intersections of a set of line segments.

Triangulation and Partitioning: Triangulation is a catchword for the more general problem of subdividing a complex domain into a disjoint collection of “simple” objects (see Fig. 2). The simplest region into which one can decompose a planar object is a triangle (a *tetrahedron* in 3-d and *simplex* in general). We will discuss how to subdivide a


polygon into triangles and later in the semester discuss more general subdivisions into trapezoids.

Optimization and Linear Programming: Many optimization problems in computational geometry can be stated in the form of *linear programming*, namely, finding the extreme points (e.g. highest or lowest) that satisfies a collection of linear inequalities. Linear programming is an important problem in the combinatorial optimization, and people often need to solve such problems in hundred to perhaps thousand dimensional spaces. However there are many interesting problems that can be posed as low dimensional linear programming problems or variants thereof. One example is computing the smallest circular disk that encloses a set of points (see Fig. 3). In low-dimensional spaces, very simple efficient solutions exist.

Voronoi Diagrams and Delaunay Triangulations: Given a set S of points in space, one of the most important problems is the nearest neighbor problem. Given a point that is not in S which point of S is closest to it? One of the techniques used for solving this problem is to subdivide space into regions, according to which point is closest. This gives rise to a geometric partition of space called a *Voronoi diagram* (see Fig. 3). This geometric structure arises in many applications of geometry. The dual structure, called a *Delaunay triangulation* also has many interesting properties.



Smallest enclosing disk



Voronoi diagram and Delaunay triangulation

Fig. 3: Voronoi diagram and Delaunay triangulation.

Line Arrangements and Duality: Perhaps one of the most important mathematical structures in computational geometry is that of an arrangement of lines (or generally the arrangement of curves and surfaces). Given n lines in the plane, an arrangement is just the graph formed by considering the intersection points as vertices and line segments joining them as edges (see Fig. 4). We will show that such a structure can be constructed in $O(n^2)$ time.

The reason that this structure is so important is that many problems involving points can be transformed into problems involving lines by a method of *point-line duality*. In the plane, this is a transformation that maps lines to points and points to lines (or generally, $(d - 1)$ -dimensional hyperplanes in dimension d to points, and vice versa). For example, suppose that you want to determine whether any three points of a planar point set are collinear. This could be determined in $O(n^3)$ time by brute-force checking of each triple. However, if the points are dualized into lines, then (as we will see later this semester) this reduces to the question of whether there is a vertex of degree greater






Fig. 4: An arrangement of lines in the plane.

than four in the arrangement.

Search: Geometric search problems are of the following general form. Given a data set (e.g. points, lines, polygons) which will not change, preprocess this data set into a data structure so that some type of query can be answered as efficiently as possible. For example, consider the following problem, called *point location*. Given a subdivision of space (e.g., a Delaunay triangulation), determine the face of the subdivision that contains a given query point. Another geometric search problem is the *nearest neighbor problem*: given a set of points, determine the point of the set that is closest to a given query point. Another example is *range searching*: given a set of points and a shape, called a range, either count or report the subset of points lie within the given region. The region may be a rectangle, disk, or polygonal shape, like a triangle.



Point location



Nearest-neighbor searching

Fig. 5: Geometric search problems. The point-location query determines the triangle containing q . The nearest-neighbor query determines the point p that is closest to q .

Approximation: In many real-world applications geometric inputs are subject to measurement error. In such cases it may not be necessary to compute results exactly, since the input data itself is not exact. Often the ability to produce an approximately correct solution leads to much simpler and faster algorithmic solutions. Consider for example the problem of computing the diameter (that is, the maximum pairwise distance) among a set of n points in space. In the plane efficient solutions are known for this problem. In higher dimensions it is quite hard to solve this problem exactly in much less than the brute-force time of $O(n^2)$. It is easy to construct input instances in which many pairs of

points are very close to the diametrical distance. Suppose however that you are willing to settle for an approximation, say a pair of points at distance at least $(1 - \varepsilon)\Delta$, where Δ is the diameter and $\varepsilon > 0$ is an approximation parameter set by the user. There exist algorithms whose running time is nearly linear in n , assuming that ε is a fixed constant. As ε approaches zero, the running time increases.

...and more: The above examples are just a small number of numerous types of problems that are considered in computational geometry. Throughout the semester we will be exploring these and many others.

Computational Model: (Optional) We should say a few words about the model of computation that we will be using throughout in this course. It is called the *real RAM*. The “real” refers to real numbers (not the realism of the model!) and RAM is an acronym for “random-access machine”, which distinguishes it from other computational models, like Turing machines, which assume memory is stored on tape. The real RAM is a mathematical model of computation in which it is possible to compute with exact real numbers instead of the binary fixed-point or floating-point numbers used by most actual computers.

Why should we care? As an example, later this semester we will study a structure called a Delaunay triangulation. The computation of this structure requires that we determine whether one point lies inside or outside the circle defined by three other points. In practice, this computation might be done with floating-point arithmetic, but floating-point round-off errors can cause the algorithm to produce the wrong result or may even cause the algorithm to fail. With the real RAM, we can assume that arithmetic operations are performed exactly, which allows us focus on the algorithm itself, rather than the messy implementation details of accurate numeric calculations.

Formally, the real RAM has a stored program, a memory consisting of an array of cells, each of which stores a single real number, and a central processing unit with a bounded number of registers. Indirection (pointers) are supported as well. (The standard RAM model of computation differs in that each memory cell stores an *integer* of arbitrary size.) The allowed operations typically include addition, subtraction, multiplication, and division, as well as comparisons. We will also allow square roots, which is useful when computing Euclidean distances. Note that operations of modulus, integer division, and rounding to integers are explicitly forbidden from the model. This is not an accident. While it is not immediately obvious, allowing these integer operations would make it possible to solve PSPACE-hard problems in polynomial time.

In spite of the unusual power of this model, it is possible to simulate this model of computation for “typical” geometric computations. This is done through the use of so-called *floating-point filters*, which dynamically determine the degree of accuracy required in order to resolve comparisons exactly. The CGAL library supports exact geometric computations through this mechanism.